

Perceptual-Recognitional Abilities and Perceptual Knowledge

Alan Millar

1. *The topic*

Disjunctivism is interesting, at least in part, because it provides a way of making sense of a relational conception of sensory experience in the face of an obvious *prima facie* objection. According to the relational conception, when I see something, for instance a book on the table in front of me, I have an experience that is *essentially relational* in that it would not be that very experience unless the book were there on the table. The *prima facie* objection, in its most straightforward form, rests on the idea that there is a perfect hallucinatory counterpart to my experience—an experience such that everything looks to me just as it does in the perceptual case, though there is no book on any table before me. The objection is that the experience in the perceptual case cannot be essentially relational since it is the very same experience as the one in the hallucinatory counterpart case—a case in which no book is present. Under disjunctivism, in the form under consideration here, it is conceded that the way it looks to the subject that things are in the perceptual case is the same as the way it looks to the subject that things are in the hallucinatory counterpart case. But it is denied that this entails that the subject would have the same experience in these cases. Consider any statement as to how it looks to the subject that things are. Under disjunctivism, that statement is such that it could be made true by the subject's having an essentially relational experience. Or it could be made true by the subject's having a perfectly hallucinatory experience. And there are other possibilities involving admixtures of relationality and partial hallucination or illusion. There need be no psychological state in common between these different situations in virtue of which it is true that it looks to the subject as if such-and-such is so.

For the purposes of the present discussion my interest in the relational conception of experience is epistemological. I take it that a major attraction of this conception is that it is thought to provide the only way to do justice to our perceptual knowledge. A guiding idea at this point is that perceptual knowledge involves cognitive contact with objects—a connection between those objects and the perceiving subject that makes it possible for the subject to have

demonstrative thoughts about, and know facts about those very objects.¹ Relationalists think that nothing short of an essentially relational experience of the book on the table can account for cognitive contact through perception. Surely, they think, it is the experience I have in seeing the book that makes it possible for me to judge and know things about *this book* as a result of seeing it. If the experience is one I could have even if no book were there, then how can it fulfil this role?

The present discussion is motivated by the thought that relationalists are right to suppose that our perceptual knowledge puts us contact with objects in a manner that enables us to demonstrative thoughts, and know things, about those objects. I present a picture of perceptual knowledge to which perceptual-recognitional abilities are central.² I argue that the exercise of these abilities in response to the occurrence of experiences is what puts us in cognitive contact with objects and facts. I then raise the question whether the exercise of these abilities depends on our having essentially relational sensory experiences. It is, of course, agreed on all sides that *seeing* the book on the table is an essentially relational matter. The issue is whether the visual experiences one has when looking at the book have to be conceived as essentially relational. On a traditional (non-relational) conception of experience the right answer is, ‘No’. That is because seeing the book is a matter, not only of having an appropriate visual experience, but also of having such an experience because of the presence of the book. I am not persuaded that this traditional answer is incorrect so far as it goes. But traditionalists have not in general provided satisfying explanations of how the judgements we are prompted to make in response to experiences amount to knowledge. As I see it an account that invokes perceptual-recognitional abilities, when suitably supplemented by an account of perceptual-discriminative abilities, can do the business, without help from a relational conception of experience.

2. Perceptual recognitional abilities: a preliminary sketch

¹ McDowell (1995/1998: 403) speaks of a ‘cognitive purchase on an objective fact’ and Scott Sturgeon (2000: 31) presents the disjunctivist’s account as one relying on the idea that veridical perception is ‘brute contact between mind and truth’. I make the metaphor of cognitive contact central to a discussion of disjunctivism in Millar (forthcoming)

Suppose I am a professional gardener and horticulturalist. Working for a client I need to find out whether or not the shrubs growing in a particular plot of a garden are azaleas. The client particularly wishes to know and is relying on me to tell her. I go to the plot and look. I see, and thereby come to know, that the shrubs are azaleas. In this situation I exercise a rather specific recognitional ability—an ability to tell by looking whether or not something before me is an azalea. This ability depends, of course, on my possessing the concept of an azalea. But it involves a lot more than that since it is an ability which I am not bound to have just in virtue of possessing the concept. It depends on my being able to recognise azaleas by sight. My possession of the concept does not so depend. (The blind can think about azaleas.)

Perceptual-recognitional abilities under this conception are not theoretical posits. We routinely think of people as having the ability to tell by looking, hearing, feeling, and so on, whether or not something is so. Our ascriptions of knowledge are often constrained by consideration of whether the subjects in question have or lack such an ability and by whether they are in a position to exercise it. Suppose that I am an electrician working with a partner. I need to know whether the bare wires I am looking are live. My partner is near the mains box. I ask him to tell me whether or not the power is off, assuming that he can find out by looking. In employing this method of checking whether the power is off I rely on knowing that my partner has the ability to tell such a thing by looking at a power box and that he is or can readily put himself in a position to exercise this ability in these circumstances.

I take it that the following points, which deserve emphasis, fall within the scope of commonsense thinking or a modest development of such thinking.

- (1) The upshot of the exercise of my ability to tell by looking whether or not the shrubs are azaleas is that I find out whether or not they are. Finding out is simply coming to know. So as a result of looking I come to know either that what I am looking at are azaleas or that they are not.
- (2) If I had judged falsely that the plants in the plot were azaleas I would not have exercised the recognitional ability in question. The general point here is that the notion of the

² Elements of the position outlined are presented in Millar (forthcoming). Here I develop the conception of perceptual-recognitional abilities somewhat further and relate it to issues about justification for belief and about why

exercise of a recognitional ability is a *success* notion. Success in relation to finding out whether or not *p* is coming to know whether or not *p*. This is not to deny that there are abilities that are abilities to do something a good proportion of the times on which one tries. My ability to hit a darts board within the 25 ring might like that. In this case success, and thus whether or not the ability has been exercised, is not to be measured by a single shot, but by the frequency of times I hit within the ring when trying to do so. The ability to tell whether or not something is an azalea is not like that. This ability is not *to judge truly enough of the time*. That ability could be exercised in making a false judgement. Rather the ability in question is an ability to tell by looking, and thus come to know, whether or not something is an azalea. It is exercised only if one ends up knowing.

- (3) In some circumstances in which I set out to tell whether or not certain shrubs are azaleas I might withhold judgement either way. For instance, the shrubs I am looking at may be shorn of leaves and flowers, and severely pruned. In this case too I would not exercise my ability to tell by looking whether or not something is an azalea, but I would exercise a closely related ability—an ability to know when to hold back from judging either way. This requires me to be sensitive to whether or not I have a good enough look at the object in question to be able to tell whether or not it is clearly recognisable as an azalea.
- (4) Even though I have the ability to tell by looking whether or something is an azalea I may on occasion judge falsely on this matter in response to my current visual experiences. I may be distracted so that I am neither sufficiently attentive to the features of the plant nor sufficiently attentive to the factors that in more propitious circumstances, or with greater care on my part, would lead me to judge the other way, or suspend judgement. The aetiology of my judgement will be in important respects like the aetiology of the judgements I make when I judge rightly in response to current visual experiences. But, to repeat, when I do judge falsely, I do not exercise the recognitional ability in question.

3. *The structure of perceptual-recognitional abilities*

knowledge matters that are not explored in the latter article.

Having a recognitional ability is in part a matter of being disposed to go through a certain judgement-forming procedure in response to suitable prompts. For instance, in the situation in which I aim to find out by looking whether or not the shrubs in the plot are azaleas, I look at them and I have certain visual experiences. If the experiences are of the right sort then, absent countervailing factors, either I judge that the thing is an azalea, or I judge that it is not an azalea.³ This latter stage of the procedure might not get started because the experiences are not of the right sort. Perhaps I don't get a good enough look at the objects in question or they don't seem to have enough of the look of azaleas. Whether the experiences lead to a judgement either way, or whether they prompt me to suspend judgement, will depend on the complex set of sensitivities that are constitutive of the ability to tell by looking whether or not something is an azalea. These will reflect the training and experience I have had, which will have endowed me with knowledge, for instance, of what azaleas look like, of other shrubs with which they might be confused, and of the range of colours they have. This background knowledge will inform my sensitivity to how good a look I need to have and how close the object has to be to the visual types characteristic of azaleas with which I am familiar, if I am to judge either way. The upshot is that my ability is not an ability to tell of absolutely anything in sight whether or not it is an azalea. Sometimes I shall not be able to tell either way.⁴

In my usual circumstances the judgement-forming procedure on which the exercise of my recognitional ability depends will yield true judgements with a high degree of reliability. But we can imagine my being in an environment in which I am prone to go through just such a judgement-forming procedure but in which I do not as a result reliably come out with true judgements. Suppose that it comes to be the case that a high proportion of what look like azaleas are replicas skilfully made from plastic and silk so that even from fairly close by an expert could not tell the difference between the replicas and the azaleas just by looking.⁵ If I am in such circumstances but unaware of the presence of the replicas then I shall go through the same judgement-forming procedure that I did in the circumstances that favoured my telling whether or not something is an azalea, but I shall not reliably make true judgements. It follows that going

³ I do not suggest that making a judgement as to whether the experiences are of the sort is part of the procedure.

⁴ This imports a degree of indeterminacy into the notion of an ability to tell by looking whether or not something is an F. People may count as having such an ability in an environment in which the only azaleas are easily recognisable as such by that person, even if they would be unable to identify azaleas in other environments in which other sorts of azaleas are common. Others may have a more finely honed ability.

through the judgement-forming procedure given suitable prompts, and other conditions, is not sufficient for me to exercise the ability to tell by looking whether or not something is an azalea. Indeed, in these circumstances I do not count as having the ability in question even though in the absence of knowledge of the presence of the replicas I am prone to go through the relevant judgement-forming procedure. I have the ability only if I am in *favourable circumstances*. Indeed, the ability is, strictly speaking, an ability to tell in favourable circumstances. The range of such circumstances is fixed by the range in which implementing the procedure results in true judgements with a high degree of reliability. Even in favourable circumstances I may go through the procedure and judge falsely because of the presence of something that looks just like an azalea but isn't. In such circumstances it could happen that a look-alike is present, though that would be highly unusual.⁶ So there is an element of luck attached to whether or not the recognitional ability in question is exercised. It is not that it is lucky that going through the procedure results in a correct judgement. Rather, whether or not I count as having exercised the recognitional ability depends on the absence of bad luck—the unusual presence of a look-alike.

On the account I have given, an ability to tell by looking whether or not an F is present is relational in this sense: it is constitutive of the exercise of such an ability that it puts one in cognitive contact with Fs and with the fact that they are Fs, or with things that are not Fs and with the fact that they are not. This contact is such that the experiences I have and the judgements I make are causally dependent on features of the objects in question. The upshot is knowing of something that is present either that it is an F or that it isn't. It is compatible with this that I could have such an ability and never have encountered Fs. Suppose, for instance, that I am undergoing training to identify certain aircraft in flight. I acquire the relevant ability by viewing realistic 3D animations of the aircraft but never see a real one. I pass the training by showing that I can come out with the right answer in response to the simulated presence of aircraft. When I do this I have not been exercising an ability to tell by looking whether or not such-and-such a type of aircraft is present. Since no such objects were present to me I have not exercised any such ability. But I have acquired an ability to tell by looking in environments in which there will be aircraft, whether or not those that I am looking at are of this or that type.

⁵ I invoke here a Henry-and-the-barn type example. See Goldman 1976.

⁶ The situation here must be distinguished from that in which replicas abound. Here the environment may be one relative to which I have the ability in question. Where replicas abound, which I cannot discriminate from the real thing by looking, I do not have the ability.

It might seem that explaining why one gains knowledge that p in terms of the exercise of a recognitional ability to tell and thus come to know whether or not p is a bit like explaining why people fall asleep in terms of their having taken a drug with the power to make someone fall asleep. Explaining how I come to know something in terms of the exercise of a capacity to come to know such things might seem no more illuminating than explaining why I fall asleep as a result of the activation of something with the power to make me fall asleep. Note first that even the *virtus dormitiva* explanation is not wholly lacking in explanatory power, since having taken something that makes one fall asleep is different from having naturally fallen asleep or having fallen asleep due to illness. But it sheds no light on what it was about the drug such that taking it makes one sleep. The explanation of particular cases of perceptual knowledge in terms of the exercise of an appropriate perceptual-recognitional ability is not like that. It is not a matter of explaining how knowledge is acquired by some power or other—we know not what—to acquire such knowledge. It is a matter of explaining, for instance, how one acquires knowledge that it's a chaffinch at the bird table through the exercise of a visual-recognitional ability. For practical purposes we have a good enough idea of the nature of this ability. It involves knowing what a chaffinch is, seeing a chaffinch, and judging from its features that it is a chaffinch. Thus it involves being responsive to the shape of the bird, its size, how it moves, and so on. All this is at the level of common sense. Reflection shows that the ability exercised on such an occasion implicates a reliable judgement-forming procedure. This further helps to make sense of why the upshot of the exercise of the ability is knowledge, since a case in which the procedure is triggered through an experience the having of which is explained by the presence of a chaffinch will be one in which the judgement made in response to the experience is true and depends on the chaffinch's being present in the way one would expect if the subject had recognised that the bird was a chaffinch. My point is that we understand the ability in question to be not just an ability to acquire the knowledge in question, though we know not how; it is an ability we know to have a certain structure, which can be described both at a common sense level and at a somewhat more sophisticated level.

4. *How experiences fit into the picture*

The picture I have presented is one on which exercising a recognitional ability is a relational matter. It puts me in cognitive contact with some object such that I know that it is or is not such-and-such. Crucially for present purposes accepting it does not commit us to adopting the view that experiences are essentially relational.

Let us imagine two scenarios—an epistemically good case and an epistemically bad case. These might be a case in which I am looking at an azalea and tell by looking that it is such, and a case in which I am looking at a skilful artificial replica of an azalea and falsely judge that it is an azalea. On the traditional conception of visual experience, the experiences I have in those cases could be the same. That poses a problem according to defenders of the relational conception. To their way of thinking if I know in the good case that must be because the experience has disclosed that the thing in question is an azalea. But if the experience in the good case is no different from the experience in the bad case then it cannot have disclosed that the thing is an azalea. The underlying thought here is that there should be an asymmetry between the good case and the bad case beyond that consisting in the fact that it is true in the good case and false in the bad case that an azalea is present. An obvious response to the problem, in keeping with the traditional conception, is to add a causal condition for cases of perceptual knowledge. On this line of thought, the asymmetry does not consist simply in the fact that in the good case the truth requirement for knowledge is met and in the bad case it is not, but in the fact that in the good case there is a *real connection* between the presence of the object and my forming a true judgement about it: my judgement depends causally on the presence of the object. Adding a causal condition also helps with Gettier counterparts of good cases in which the truth requirement is met, for instance, cases in which a life-size photograph interposes between the object in question and the subject. In the Gettier cases there is no real connection between the presence of object and the formation of the judgement, and so no genuine cognitive contact with the objects in question. Yet the mere addition of causal condition falls short of providing a satisfying explanation of why knowledge is acquired only when it is satisfied. For one thing, the fact that a causal condition would block counterexamples to knowing does not amount to the needed explanation. We need to understand is what it is about perceptual knowledge that demands that the causal condition be satisfied. For another, the form that such a causal condition

should take is notoriously unclear because not every causal dependence that judgements have on experiences and experiences have on the presence of objects will do the trick. We need to better understand what constrains the causal dependencies that count.

Defenders of the relational conception of experience circumvent these problems. Instead of thinking of the experiences gained in perception of some object as being causally dependent on the presence of the object, they think of the object as being a constituent of the experience. We are to take it that the experience puts us in touch with an azalea in the good case, just because the azalea seen is an essential ingredient of the experience in the good case—it would not have been that very experience if that very azalea had not been present. So it is on that account not the same experience as the experience in the good case.

A satisfying account of the asymmetries must explain how it is that something that the subject has done in the good case, that was not done in the bad case, or in Gettier counterpart cases, secures knowledge. This is where the epistemologically minded defender of the relational conception of perceptual experience gets a foothold, pushing the claim that the achievement in the good case, and only in the good case, consists in one's having based one's judgement on an experience of which the object in question is a constituent. My own view is that to explain the asymmetries adequately we do not need to assume a relational conception of experience. Returning to my examples, the crucial asymmetry is that in the good case I have exercised the ability to tell by looking whether or not something is an azalea and in the bad case, and in the Gettier counterpart case, I have not. Note that in the good case there is a relation of causal dependence between the experiences that trigger the judgement and the presence of the object. But not just any causal dependence will do. It must be a dependence that is compatible with the whole process being an exercise of the relevant recognitional ability.

On the account I have provided then perceptual knowledge is explained in terms of the exercise of recognitional abilities. It is conceded that the experiences gained in perception do not suffice to account for our being in cognitive contact with the objects we see. Rather it is experiences gained in circumstances in which I have exercised the relevant recognitional ability. The experiences matter because they played a role in making it the case that the ability has been exercised.

The account needs supplementation if it is to provide a satisfactory response to defender of the relational conception. When I tell that an object is an azalea by looking I will, of course, have

picked out the object by sight. What is it to do that? What is it that makes the object available to me as something about which I can go on to make a judgement as to whether it is or is not an azalea?⁷ Do we not need to assume a relational conception of experience to make sense of that? I think not, but we do need to supplement the story of perceptual-recognitional abilities with a story about perceptual-discriminatory abilities. On a traditional conception of experience we cannot account for visually picking out an object simply in terms of having a visual experience of an appropriate (non-relational) character since such an experience can be had when no object is picked out. Nor can we account for it by adding a causal requirement, since simply adding a causal requirement does nothing to explain how the subject has achieved something in picking out an object by sight or which sorts of causal dependence matter. Picking out an object by sight is visually discriminating it from its background. Such discrimination, I take it, is sub-doxastic in that its immediate upshot, in addition to the formation of visual experiences, is not belief or judgement, but being primed to respond behaviourally in ways that are appropriate given what one is doing. When I pick out the azalea I not only have certain experiences but I am primed so that, for instance, should I wish to look closer I'd move towards the location of the shrub. When I pick out people coming towards me in a corridor I am primed so that I shall make appropriate adjustments to avoid bumping into them. Visual discrimination makes objects available for thought, and available to be recognised as objects of this or that sort, in the presence of suitable recognitional abilities. So relationality is vital to account for perceptual knowledge. There is the relationality of visual discrimination and there is the relationality of the exercise of perceptual-recognitional abilities. So far as I can see none of this requires us to adopt a relational conception of experience. As I stressed earlier, it is agreed on all sides that perception is relational. The issue has been whether the experiences implicated in perception are essentially relational.

5. Comparison with traditional reliabilism

Traditional reliabilist accounts of perceptual knowledge have it that perceptual knowledge is true belief formed through the implementation of a reliable belief-forming procedure. In making reliability central such accounts are to that extent correct. But my account differs from traditional reliabilism in a number of respects.

⁷ This question is pressed by John Campbell (2002).

(i) Traditional reliabilism is an attempt to explain what knowledge is in terms of true belief plus something else. I make no such attempt. I explain particular instances of perceptual knowledge in terms of the exercise of perceptual-recognitional abilities, but the abilities are characterised epistemically—they are abilities to find out—that is, come to know—by perceiving.⁸

(ii) Under traditional reliabilism the implementation of a reliable procedure is not a success notion. The procedure might be implemented and the subject not acquire knowledge; something that is not an F could produce the kind of experiences that, by the procedure, result in a judgement that an F is present. As I have stressed already the exercise of a perceptual-recognitional ability *is* a success matter.

(iii) Traditional reliabilism is problematic if presented as a partial analysis of our ordinary concept of perceptual knowledge. Such an approach suggests a picture of how we ascribe perceptual knowledge on which such ascriptions are ultimately grounded in evidence that the subject has a true belief which results from the implementation of a reliable procedure.⁹ One problem with this is that ordinary thinking about perceptual knowledge does not supply us with a determinate conception of the procedures on which our perceptual recognitional abilities depend. Reflection shows that they must take us from experiences or certain sorts to beliefs of a corresponding sort. But we have no idea how to specify with any precision which experiences, against which background of beliefs and sensitivities, are implicated. Despite this, we are often able to tell when people know by looking that an F is present. We do so, for instance, when we see that an F is present and are able to tell from a subject's reactions that he or she has noticed the F and recognised it as an F. Suppose that I am out with a friend walking his dog. The dog has disappeared among some bushes but emerges. My friend calls out, not in the way he would if trying to locate it, but in the way he would if he had just seen it. Seeing both the dog and him I take in that he is calling to it, not calling out for it. In this situation my judgement that my friend

⁸ I am sympathetic to those who are sceptical about the possibility of analysing knowledge into true belief plus something else. See, for instance, McDowell (1993 and 1995) and Williamson (2000).

⁹ I am assuming here that an adequate conceptual analysis would reflect the constraints which in practice govern our dealings with the concept in question. So it would be a problem for such analysis if its import for how we are in practice governed were psychologically unrealistic.

has recognised the dog is not based on any evidence to the effect that, by and large, he forms true beliefs that his dog is present when his dog is in clear view. It is based (a) on the knowledge that my friend knows his dog, and can recognise it by sight, and (b) on my knowledge that on this occasion he has just noticed and recognised it by sight. My knowledge that my friend knows his dog might be based simply on seeing the two interact on this occasion. This and my knowledge that he has picked out and recognised his dog by sight are arguably the outcomes of the exercise of complex recognitional abilities, rather than conclusions drawn from prior assumptions about the present situation, viewed in the light of my friend's history of acquiring true beliefs as to the presence of his dog. In both cases I tell from visually manifest clues. For instance, I tell from the layout of the environment and from the orientation and behaviour of my friend that he has seen and recognised his dog. Telling from these factors should not be conceived as a matter of inferring from assumptions as to the layout and about the behaviour and orientation of my friend. I would be hard put to capture by description the clues to his having spotted the dog, just as I would be hard put to capture by description the features of a familiar face on the basis of which I recognise whose face it is. The model we should work with is recognition rather reasoning from evidence. This helps to explain how my engagement with my friend's cognitive state is at the level of his visual-cognitive contact with the dog.¹⁰ That is to say, it is the level at which success notions like seeing and recognising are in play. Reflection on the situation in which the seeing and recognising has been achieved brings into play the idea that a reliable judgement-forming procedure must have been implemented. But the order of knowledge is from recognition that there has been seeing to the conclusion that a reliable judgement-forming procedure has been implemented, rather than vice versa. (I discuss related points below in section 8.)

Recognitional abilities, as I conceive them, bear some resemblance to what are sometimes called *intellectual virtues*. For Ernest Sosa, to have an intellectual virtue relative to a certain environment, with respect to a field of propositions and set of conditions, is to be so constituted that if one is in that environment, in the relevant conditions, and believes or disbelieves some proposition within the relevant field, then one will very likely be right (Sosa 1991: 284). Like recognitional abilities, virtues in this sense are tied to specific subject-matters, and one counts as

¹⁰ An implication of this view is that we need a liberal account of what perceptual knowledge can embrace. See my (2000).

having them only if one is in a suitable environment. But the exercise of a virtue is not a success notion. One can exercise the virtue in the right environment and in the right conditions and believe or disbelieve wrongly.

There is an obvious kinship between my account of perceptual recognitional capacities and Alvin Goldman's well-known relevant alternatives account of perceptual knowledge.¹¹ Goldman takes us to have perceptual knowledge that *p* if, roughly, we are caused to believe that *p* non-inferentially through have an appropriate experience in circumstances in which there is no relevant alternative to its being the case that *p*. My own account captures what I take to be a key insight contained within this account. It is the actual circumstances that determine what relevant alternatives are. If the actual circumstances do not throw up things that are not *F*s but which look just like them, or do so only rarely, then a subject can know by looking that an *F* is present without determining that what is present is not a look-alike. However, the full definition of perceptual knowledge that Goldman provides is so complex that it is hard to see how it could govern our actual applications of the concept of knowledge. From the standpoint taken here that is a defect. We want an account of knowledge to enable us to understand the role of knowledge and of ascriptions of knowledge in our social interactions.

I do not think there is much doubt that our most natural ways of thinking about perceptual knowledge implicate recognitional abilities as I conceive them. But why should that be? Why should success notions be so crucial to our thinking? At least part of the answer has to do with why knowledge matters to us. I address this below in section 7. Before that I consider how perceptual knowledge as understood here links up with notions like that of justified belief. This is important because without some discussion along those lines the suggested approach to perceptual knowledge might seem to avoid the issues that have preoccupied traditional epistemologists.

6. *Justification*

When I look at the shrubs and tell that they are azaleas I come to be in a position such that I am justified in being sure that they are azaleas. How should we think of this justification? Should we

think of it as explaining how I know or should we think of it as deriving from the fact that I know? The view I take here is that the latter is the right way to think of the matter. The explanation for why I know is simply that I have exercised a way of telling—specifically, a perceptual-recognitional ability to tell whether or not something is or is not an azalea from the way it looks. That said, at least in typical situations when we know by looking that something is an F we are aware, that is know, that we have seen the thing to be an F. In those situations the visual experiences that, via the exercise of the relevant perceptual recognitional ability, enable us to know that something is an F, also enable us to tell that we see that thing to be an F. This latter knowledge is made possible by the exercise of a higher-order recognitional ability—an ability to tell whether or not one sees that an F is there. This latter ability is not an introspective ability, for there is no inner scrutiny of anything. Just as I have learned to tell when something is an azalia in response to certain visual experiences, so I have learned to tell when I see an azalea in response to *the very same experiences*. I do not think about or introspect my experiences in the case in which I judge that I see azaleas any more than I do when I tell that the shrubs are azaleas by way of response to the same experiences. I go through a procedure that takes me from experiences to a judgement and a belief that I see that the shrubs are azaleas. This procedure is constitutive of an ability to tell, in a suitable environment and under suitable conditions, whether or not I see that certain things are azaleas. The upshot of the exercise of such an ability is that I become apprised of a relational fact about myself and about the azaleas, and not as a result of any inference.¹²

Let us suppose then that I am looking at the azaleas and recognise them for what they are. I am sure that they are azaleas. What justifies this assurance? A natural answer is that what justifies the assurance is my knowing that I see that the shrubs are azaleas. Where I to explain why I know that the shrubs are azaleas it would be natural to cite the fact that I can see that they are. But it is also true that were I to justify thinking that they are azaleas it would be natural to do so in terms of this same fact.

¹¹ See Goldman 1976. Reflection on the case of Henry and the barn, which is so central to Goldman's account, undoubtedly contributed to way I have represented recognitional abilities and their relation to reliable judgement-forming procedures.

¹² I have reflective access, as one might say, to that in virtue of which I know that the shrub is an azalea, but it is not restricted to modes of access—introspective and a priori—recognised under standard internalist theories of knowledge or justification.

The view I am presenting links being justified with having reasons in an entirely natural way.¹³ It is because I know that I see that the shrubs are azaleas that the consideration that I see that they are azaleas can serve as a reason I have to believe that they are azaleas. And because I will retain in memory that I have seen that the shrubs are azaleas, the consideration that I have seen that they are can serve as a reason to (continue to) believe that they are.

The preceding remarks naturally lead us to consider whether I am guaranteed to have a good reason to believe that the shrubs are azaleas simply through having exercised the relevant recognitional ability. Is it built into my having exercised the ability that I will know that I have seen that the shrubs are azaleas and on that account be justified in believing that they are? There is no question that at least typically I would know both that the shrubs are azaleas and that I see that they are, and on that account be justified in believing that they are. The view I am presenting does not commit us to supposing that perceptual-recognitional knowledge that *p* guarantees that one is justified in believing. But it is important for the use to which we put the concept of knowledge that we can in perceptual cases readily tell that we know. I pursue this further in the next section.

More traditional approaches to the theory of knowledge accommodate the idea that one's reason for believing that *p* might be that one sees that *p*, but the availability of the consideration that one sees that *p* has to be explained in terms reasons for thinking that one sees that do not entail that one sees. The natural way to do this is in terms of the consideration that it is to one as if one sees that *p*, or some such thing. The question then arises about what justifies one's thinking that its being to one as if one sees that *p* is the upshot of its being the case that *p*. I suspect that what has made such approaches seem compelling is the thought that justification for the application of success notions like seeing that *p* cannot derive from simply recognising that one sees that *p*. It is crucial for the view I am presenting that such recognition is commonplace.

7. Why knowing matters

My focus on knowing through the exercise of perceptual-recognitional ways of telling fits neatly with a plausible conception of why knowing matters in a wide range of social interactions. It can matter to me that I know of the shrubs I am looking at that they are azaleas just because I am

¹³ Compare McDowell 1993 at this point.

called upon to vouch for whether or not they are by my client, who wants to know from me. Her confidence in what I tell her is grounded in confidence that I can find out and have done so. My confidence in what I vouch for is grounded in my confidence that I have seen what the shrubs are.

Knowledge enables one to have well-grounded assurance. We can have well-grounded assurance that *p* in virtue of the fact that we see that *p* or have seen that *p*, and on that account are justified in believing that *p*. One reason why we value well-grounded assurance is that it enables us legitimately to vouch for something's being so. When we vouch for something's being so we give it to be understood that we have the kind of assurance that is grounded in knowledge. I do not deny that we often say what we think is true when we do not know, and on reflection would accept that we do not know. I do not regard such saying as vouchings. We vouch for something's being so in contexts in which someone wants to know whether it is and we wish to oblige them. It is part of the practice of vouching that you should not vouch unless you know.

Another reason why we value the assurance that is grounded in knowledge is that we need our true beliefs not merely to be true but to be secure (not easily vulnerable) as well. We do not want our true beliefs to be like the statues of Daedulus, which slip away when untethered. (*Meno*, 97D-E). Suppose that I believe correctly that I have sent off a letter of recommendation for a student, but I do so, not because I remember sending it, but on the grounds that a colleague has told me that I told him that I had sent it just after doing so. But my colleague is confused. *I* did not tell him I had sent off the letter. Someone else had told him that *he* (this someone else) had sent off a letter of recommendation. When my colleague spoke to me he thought it was I who had told him. As things stand my only basis for thinking that I have sent off the letter is his false report. So my belief is liable to be given up if I learn that my colleague was confused, for then I would realise that I don't know. But if I distinctly remember putting the letter in the mail tray it would take a lot more than this to dislodge my belief that I have sent it.¹⁴

It can be practically important that we not give way to doubt with respect to a true belief. My true belief that I sent the letter was easily abandoned in the face of the realization that my colleague's report was false. I then had to make time-consuming enquiries to settle whether I did

¹⁴ The example is similar to Williamson's burglar example (2000: 62). I think it is important that the knowledge plays its characteristic explanatory role in virtue of the subject's recognition that he knows in a certain manner.

or did not send it. Of course, it could be that a true belief that *p* that is not knowledge that *p* is not abandoned because the factor that prevents the subject from having knowledge that *p* never comes to light. That does not tell against the importance of knowledge as opposed to merely true belief. Although merely true belief *might* do as well as knowledge, it is liable not to because of the possibility that the knowledge-preventing factor does emerge.

Knowledge is instrumentally valuable then, not just because the implicated true beliefs are instrumentally valuable in virtue of being true. The instrumental value of knowledge, over and above the instrumental value accruing to the implicated true beliefs in virtue of their being true, lies in the value of having well-grounded assurance that renders our beliefs secure.

These considerations about why knowledge matters help to explain why success notions like seeing-that are so basic to our thinking about knowledge. We want knowledge when we wish to be assured. We come by assurance as readily as we do in practical affairs because we can so readily tell when we, or others, know something in a wide range of matters that are relevant to our practical affairs. My seeing that the shrubs are azaleas provides me with well-grounded assurance that I know and can thus vouch for the shrubs' being azaleas. I can *readily* tell that I see and thus can readily tell that I know. As I have already suggested, the very same experiences that figure in my knowledge of the first-order worldly fact that the shrubs are azaleas, can also engage a higher-order recognitional ability, resulting in my knowing the relational fact that I see that the shrubs are azaleas. My client's assurance that the shrubs are azaleas derives from her knowledge that I can tell in such matters by looking and have done so on this occasion.

If the concept of knowledge were governed by a more or less complex definition in terms of necessary and sufficient conditions then knowledge would be more like a theoretical posit than it is, at least when perceptual knowledge is at issue. Applying the concept knowledgeably would require that one should know that the conditions for its application have been satisfied. Since it would not be at all clear that we have ready access to whether or not the conditions are satisfied it would be harder to understand how knowledge can play the role it does in interactions such as I have described.

A related point is that it is important to distinguish between, on the one hand, assurance that *p* deriving from personally seeing that *p* or knowing that someone else was in a position to see that *p* and a high degree of confidence that *p* deriving from, say, inductive evidence. I may have a high degree of confidence that if I go to a certain café at a certain time I shall see some friends

there at that time. My confidence may be grounded in the knowledge that they not only go there regularly at that time but that it is something of a ritual for them to do so. But despite the ritual character of their behaviour their attendance is not perfect; things can crop up that prevent their going. So I do not take myself to *know* that they will be there. I might tell others that it's a good bet that they will be there but I would not vouch for their being there. But if I am at the café with them I could vouch for their being there to any interlocutor with whom I happened to be in telephone conversation. In this latter situation it is not merely that the level of my confidence is sufficiently ratcheted up to make it reasonable to vouch. There is a difference in kind and not merely degree between how things stand with me when I know through seeing something is so and how things stand when I have good evidence that makes it very likely that something is so. So it would be odd if knowledge could be had in virtue of having a true belief based on evidence that merely made it more or less probable that p.¹⁵ Given that available evidence so often does no more than make it more or less probable that something is so, it might be tempting to think that we would be better not bothering about knowledge and settle for probability. For large tracts of our beliefs this might well be the right approach to take. But our social interactions require us to discriminate between situations in which we, or others, can vouch for something's being so and those in which we, or they, are at best in a position to be reasonably confident.

The foregoing takes us to a further point of contrast with traditional reliabilism about knowledge. Consider a case in which we do not have independent access to whether or not p, but wish to tell whether or not somebody knows that p by looking. In such a case the ascription of knowledge under reliabilism would be based on the assumption that a reliable procedure has been implemented and the chance that the belief is true is, therefore, high. That is not how in practice we think of knowledge in perceptual cases.¹⁶ The typical situation in which we ascribe perceptual knowledge that p to others, in circumstances in which we do not ourselves know that p, is one in which we rely on the following assumptions: (i) that they are in a position to tell whether or not p by perceiving, and (ii) that what they tell us on the matter in this situation will

¹⁵ This is an important strand in the thinking of McDowell 1982. It is picked and developed in Travis 2005, who traces it back to Cook Wilson.

¹⁶ The same made be said of cases in which knowledge that p is based on some indicator that p—a fact that would not obtain were it not that p. To make this plausible would take us too far from the current topic. It requires that we extend the conception of a recognitional ability to embrace cases in which a fact that p is recognised to obtain through seeing that something else is so. A key question would be how to makes sense of this without imputing to

be true because based on what they have perceived to be so. When they tell us, say, that *p*, we are in a position to infer that they know that *p*. We reach this conclusion not on the basis of assumptions that if true would raise the chance that what they tell us is true. Of course, if this picture is to contrast with one in which a hazardous leap is made to the conclusion that the interlocutor knows, from something that makes it highly probable that the subject knows, the various assumptions involved should not themselves result from hazardous leaps from evidence. I think that this condition can be met, though the matter demands much fuller treatment than I can give here. Still, it seems fairly clear that we can know that they are in a position to tell that whether or not *p* by looking because we know that they are located so as to be in such a position (for instance, beside the power box with the mains switch). And we can know that they would not vouch for its being the case that *p* unless they had perceived that *p* because we know them and can recognise when the situation is one in which what they vouch for is to be relied upon (for instance, whether or not the switch is off).

My claim is that perceptual knowledge can play the role that it does in social interaction because that there is a basic level of engagement with the perceptual knowledge of others at which we can tell that they know that something is so as a result of a chain of knowledge-transmission of the sort I have described. At this level we think of subjects as being or having been in perceptual cognitive contact with the fact. We do not think of their coming to believe that *p* as resulting from the operation of a reliable process that merely raising to a high degree the chance that it's a fact that *p* and then make a further leap to the conclusion that they know. This is not to deny that in situations in which subjects tell that *p* because they see that *p* a reliable, yet fallible, judgement-forming procedure will have been implemented. On the account of perceptual-recognitional abilities proposed here such a procedure will have been implemented when a subject sees that *p*. The point, to repeat, has to do with the level at which we conceptualise what is going on when knowledge is had in perceptual cases.¹⁷

the subject an inductively justified belief in a covering generalisation. Closely related themes are pursued in Travis 2005. See Millar 2005 for some critical reaction to Travis's discussion.

¹⁷ Work for this article is part of project on the Value of Knowledge funded by the Arts and Humanities Research Council. I am grateful to the Council for its support. For numerous helpful discussions am also grateful to Duncan Pritchard, and Adrian Haddock, who are key participants in the Value of Knowledge project, and to the audience at the Glasgow Disjunctivism Conference in September 2005 for discussion of an earlier version of this article.

References

- Campbell, John (2002) *Reference and Consciousness* (Oxford: Clarendon Press).
- Goldman, Alvin (1976) 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73, 771-91.
- McDowell, John (1993) 'Knowledge by Hearsay' in B. K. Matilal and A. Chakrabarti (eds.) *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony* (Dordrecht: Kluwer). Reprinted in McDowell 1998, 195-224.
- McDowell, John (1995) 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93. Reprinted in McDowell 1998.
- McDowell, John (1998) *Meaning, Knowledge and Reality*. (Cambridge, MA.: Harvard University Press).
- Millar, Alan (2000) 'The Scope of Perceptual Knowledge', *Philosophy* 75, 73-88.
- Millar, Alan (2005) 'Travis's Sense of Occasion' *Philosophical Quarterly* 55, 337-42.
- Millar, Alan (forthcoming) 'What the Disjunctivist is Right About' *Philosophy and Phenomenological Research*.
- Sosa, Ernest (1991) *Knowledge in Perspective* (Cambridge: Cambridge University Press).
- Sturgeon, Scott (2000) *Matters of Mind: Consciousness, reason and nature* (London: Routledge).
- Travis, Charles (2005) 'A Sense of Occasion', *Philosophical Quarterly* 55, 286-314.
- Williamson, Timothy (2000) *Knowledge and Its Limits* (Oxford: Oxford University Press).